



United States Department of Agriculture

**Office of the Chief Information Officer
Network Engineering Division**

Telecommunications Enterprise Network Design

Development of Network Capacity Design

Task IV Report

March 13, 1998

Executive Summary

The purpose of this study is to make cost reduction recommendations to improve performance and utilization – without changing topology or impacting mission critical programs – of the existing USDA Data Networks. Recommendations take the form of reducing bandwidth for underutilized links and increasing bandwidth for over-utilized links – recapacitation.

The status of the Telecommunications Enterprise Network Design project is depicted graphically in the EN Development Task Flow following the Executive Summary. The primary goal of the project is to optimize the current USDA Data Networks to minimize cost while maintaining or improving performance. The optimization process is composed of several phases. Following a comprehensive network baseline analysis, a description of all network elements, the first optimization phase recommends correct link sizes based on "Average" and "Peak" traffic. The link size analysis, recapacitation, does not involve any changes to network topology. Subsequent optimization phases include rerouting of links, concentration of nodes and the identification of a network backbone.

Key Concepts

- Based on average and peak utilization information 100% of the measured USDA Data Network links are, to some degree, incorrectly sized.
- Approximately 84% of the measured USDA Data Network links have an average usage at less than 15% of the theoretical link capacity and represent links on which cost can be reduced.
- Approximately 16% of the existing USDA Data Network links have an average usage greater than 50% of theoretical capacity and require timely attention to avoid an impact on mission performance. (4% are used at >70%, requiring immediate attention.)
- Downgrading (reducing bandwidth) of underutilized network links can potentially reduce USDA service cost by as much as 23%.
- Upgrading (increasing bandwidth) of over-utilized network links can improve service performance for several USDA Agencies.

TELECOMMUNICATIONS ENTERPRISE NETWORK DESIGN

Development of Network Capacity Design – Task IV

Network Capacity Analysis

The results of the existing network link capacity evaluation demonstrate that **100% of the measured links are, to some degree, incorrectly sized.** However, only 16% of the links require timely attention to avoid an impact on mission performance. Most of the links (84%) are underutilized (<15% of theoretical capacity used). (Fig. A) Underutilized links are Cost Sensitive representing potential cost savings but do not require immediate attention since mission performance is not affected.

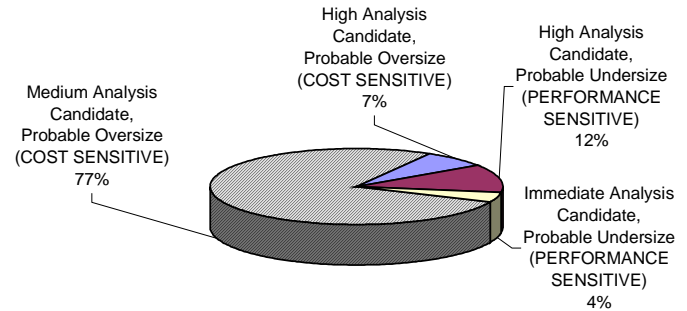
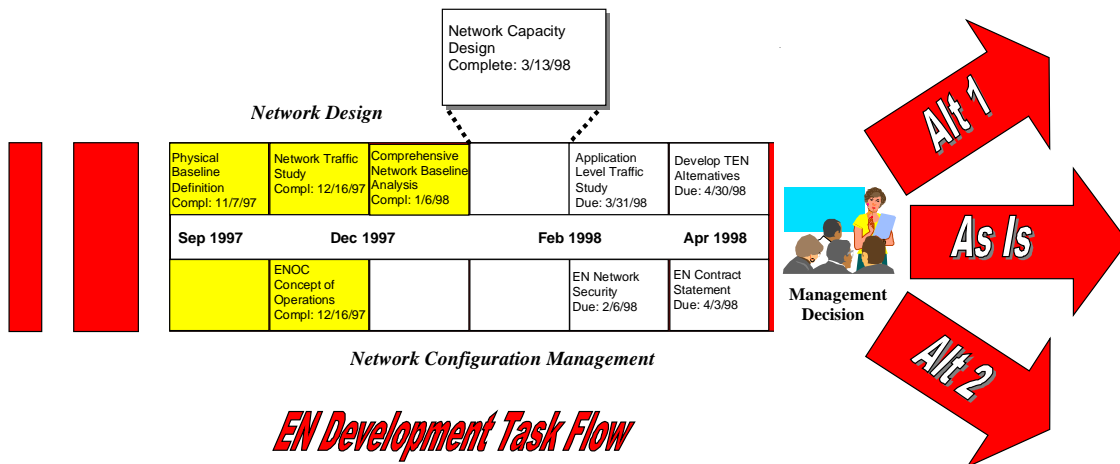


Figure A USDA Network Link Capacities

Cost Savings from Link Recapacitation

The Wide Area Network access cost of the existing USDA Data Network links and the re-sized links have been calculated based on the number of links times the access cost for each type of link. The USDA has the potential of reducing Telecommunications Network costs by as much as 23% if recapacitation recommendations are implemented.



1.0 Introduction

1.1 Objective

The Network Capacity Study provides recommendations to significantly reduce costs and to improve performance and utilization – without changing topology or impacting mission critical operations – of the existing USDA Data Networks. Recommendations are based on:

- collecting network traffic and link data describing network utilization;
- developing network link sizing metrics and link optimization strategy;
- applying link sizing metrics and link optimization strategy.

Recommendations take the form of reducing bandwidth for underutilized links and increasing bandwidth for over-utilized links.

1.2 Background

1.2.1 Goals of the Enterprise Network Project

The USDA Information Systems Technology Architecture (ISTA) describes the guidelines for the USDA to achieve the objectives of the Government Performance and Results Act (GPRA) of 1993. The Telecommunications Architecture, Part III of the ISTA, provides the framework for managing the efficient use and continued evolution of telecommunications services, systems, and networks in performing the Department's missions in a cost effective manner.

Development of a USDA Telecommunications Enterprise Network (TEN) is the cornerstone of the Department's strategic information systems plan. The TEN must also conform to the goals of the Telecommunications Network Stabilization and Migration Program (TNSMP) (Aug, 1997). The TNSMP goals are:

- **Enhance interoperability and user access** through standardization of agency networks, Internet access, use of open and de facto standards, Internet Protocol (IP) Address management, defined resource management responsibilities, international connectivity.
- **Improve program mission performance.**
- **Reduce cost through consolidation.**

1.2.2 Status of Telecommunications Enterprise Network Design Project

Using the Geographic Network Analysis Process (GNAP) *USDA Network Design Process, version 1.1, August 1997*, Network Engineering has developed a comprehensive description of the current USDA Data Networks. This description is presented in *Telecommunications Enterprise Network Design, Comprehensive Baseline Analysis: Task III, January 6, 1998, Network Engineering*. The parameters of this analysis are network equipment, service, utilization, performance, cost, and survivability. The NetMaker XA[®] System has been used to electronically discover the nature of the network hardware elements, measure network usage and performance, and calculate network costs. **In summary, the analysis found that the state of the current USDA Data Networks provides many opportunities to reduce cost while maintaining or improving performance.**

In addition to TEN design activities, the Configuration Management Team has begun to describe the framework of Enterprise Network Operations Center (ENOC). *Telecommunications Enterprise Network Configuration Management, Enterprise Network Operations Center: Concept of Operations, December 17, 1997*, defines critical Enterprise Network Operations Center activities, responsibilities, and the management interaction between the USDA and private contractors. Definition of a TEN Security System Architecture is in progress.

1.2.3 Network Level Traffic Study

The Network Level Traffic Study, Task II of the Initial TEN Design Plan, presented an analysis of USDA Network utilization based on limited data collection. **The study serves primarily as a proof-of-principle for the electronic survey process using NetMaker XA[®].** The data collected has been used to describe "Average" and "High" utilization days. The results reflect generalized network information. For Task II, the following factors did not permit determination of network utilization for specific links:

- Link speed values were limited by accuracy of router configuration.
- Data was restricted to limited information collected for the Physical Baseline Study.
- Data represented a composite of two separate collection activities.

Caveats imposed on Task II data have been resolved in subsequent network utilization data collection activities. The qualified agency specific information disseminated in Task II provided agency

representatives with network descriptions for review and correction. The response to this information has been supportive of the TEN project activities and has enabled more accurate "link specific" data collection. The improved network information has provided the details necessary to make network capacity design recommendations.

2.0 Methodology

2.1 Equipment

Network level traffic analysis is a key process used by Network Engineering to define network attributes. Traffic analysis – link capacity, link utilization, type of applications – provides the data necessary to make network performance and cost recommendations. The combined USDA Data Networks represent a very large and complex network, the accurate and timely analysis of which requires the use of electronic equipment.

The NetMaker XA^{®1} system includes the tools necessary to acquire and process network traffic information. The NetMaker XA[®] Baseline Wide Area Network (WAN) tool for the NetMaker XA[®] Interpreter tool dynamically polls network routers, using Simple Network Management Protocol (SNMP), to collect Management Information Base (MIB) statistics. For a defined time period, the polling procedure collected router statistics on the amount of traffic at device serial interfaces. The resulting information represents usage trends that identify traffic volume. Along with knowledge of link type and the applications in use, the NetMaker XA[®] Designer tool (Designer) is used to identify and to correctly categorize links as under or over-utilized links. Designer recommends under-utilized links for bandwidth reduction – maintaining performance while reducing telecommunications costs. Conversely, Designer recommends over-utilized links for bandwidth increase – improving performance at an appropriate telecommunications cost increase.

2.2 Data Collection Procedure

Network traffic is the data transferred, by whatever technology, between two WAN nodes. WAN links, in addition to routers, are fundamental building blocks for USDA Data Networks. Using several different technologies [X.25, Dedicated Transmission Service (DTS), Frame Relay (FR), Asynchronous Transmission Mode (ATM)], WAN links interconnect geographically distant LANs (Local Area Networks). The actual traffic information captured from a router includes the time/date stamp for the transmission start time and end time, the number of bytes transferred and received, the number of packets transferred and received, the traffic source and destination.

¹ NetMaker XA[®] is a registered trademark of Make Systems, Inc.

Traffic data used in the evaluation of the USDA's network utilization was captured during standard business hours normalized for the entire United States. Business hours were defined as 8:00 AM to 5:00 PM local time. To insure capture of business traffic, data was captured from 4:00 AM to 7:00 PM Mountain Standard Time. Traffic captured for this evaluation was collected weekdays from January 12, 1998 through February 6, 1998.

The NetMaker XA[®] Baseline tool was programmed to capture hourly traffic statistics for each individual USDA router from 4:00 AM to 7:00 PM (MST). These baseline traffic files were downloaded and further processed to reduce the number of traffic demands loaded into NetMaker XA[®]. Due to the size of USDA's modeled network traffic demands captured hourly over the four week time period exceed the capacity of the current NetMaker XA[®]. To mitigate this limitation, the hourly data was summarized into an hourly average, and maximum traffic demand and subsequently used for network analysis. The hourly maximum is referred to as the High traffic demand used to calculate the High Utilization value. The average hourly traffic demand was used to calculate the Average Utilization value.

Link capacities reported in this document are based on the information collected in the January 30, 1998 Discovery. Since network traffic information is read from a router MIB, the accuracy of collected information is a function of the accuracy of router configuration. If a router is incorrectly or incompletely configured, a default link speed is collected, resulting in an incorrect calculation of link utilization.

2.3 Link Sizing Metric

Two categories of logical WAN connections that USDA now uses are DTS (non-switched point-to-point transmission; private line) and FR Permanent Virtual Circuit (PVC). The size of these connections are measured in kilobits per second and the maximum number of bits is assumed to be the circuit rate for DTS and the access rate for FR.

2.3.1 Parameters for Establishing Link Utilization Metric

Measurement of link utilization and the development of a metric to use for properly sizing links is dependent on several parameters:

- Type of WAN link (DTS or FR)
- The available bandwidth for a WAN link as measured by its circuit rate or FR access rate
- The number of bits transferred over a given circuit captured during a representative period of time. This traffic will be presented as a peak load and as an average load. The peak defines the busy hour and busy day for the link

A generally accepted WAN utilization measure is the amount of data transferred divided by the available bandwidth and presented as a percentage.

$$\text{Utilization Metric} = \frac{\text{data transferred (bits/sec)}}{\text{available bandwidth (bits/sec)}} \times 100$$

Applying the utilization metric to network links provides link descriptions based on the urgency of their need for optimization. In addition to calculation of a utilization metric, link analysis for optimization requires:

- current application performance requirements,
- future application performance requirements,
- confirmation of the correct available bandwidth, and
- recommendation for re-sizing the link.

2.3.2 Circuit Analysis Priority

Present optimization efforts are primarily concerned with maintaining or improving link performance without modifying network topology. As a consequence, peak traffic utilization significantly affects the measure of circuit usage. However, a non-sustained peak usage skews upward the circuit utilization measure. To compensate for the peak usage effect, link utilization measurement must also reflect average usage. The exact nature of this peak-to-average performance relationship is somewhat intuitive and requires further study for more exact quantification.

To maintain performance consistency, link utilization should not exceed 70% of theoretical capacity. Similarly a circuit utilized less than 15% of theoretical capacity indicates a circuit that can be reduced without loss of performance. Using these values as rules of thumb, Table 1 describes the relationship between utilization (Average and High) and Analysis Priority.

Analysis Priority Categories Based on Circuit Utilization				
Average Utilization	High Utilization			
	<1%	1-15%	15-50%	>50%
0-20%	High	Medium	Medium	High
20-50%			Low	Immediate
>50%				Immediate

Table 1 Analysis Priority Categories

Definition of Analysis Priority Categories

Immediate Priority indicates the link is probably incorrectly sized for current applications and performance may be adversely affected (Performance Sensitive). The size of links may need upgrading.

High Priority indicates the link may be incorrectly sized for current applications; either oversized or undersized. If a link is undersized, it is Performance Sensitive meaning that performance may be adversely affected. Oversized links represent a potential for cost saving (Cost Sensitive).

Medium Priority indicates the link may be incorrectly sized for current applications; probably oversized. This category is Cost Sensitive and does not impact on mission performance.

Low Priority indicates the link is probably correctly sized.

2.4 Link Optimization

2.4.1 Strategic Phases

Network optimization mandated by the ISTA is defined as providing the same or better level of service for a network customer (USDA Agency) at the same or lower price. Network optimization is a multi-phase process consisting of the following activities:

- Recapacitation of network links
- Rerouting of linkages between routed nodes
- Concentration of nodes
- Identification of a network backbone

This report deals with the first phase of optimization, maintaining the linkages between routed nodes but changing the link sizes according to the determined traffic demands. For various reasons, including time constraints and equipment limitations, a complete traffic sample representing actual workload has not been measured. Current data represents an arbitrary target to make all circuits utilized to 50% of their theoretical capacity.

The culmination of network optimization is network design alternatives for an initial TEN. The presentation of the initial TEN alternatives is the objective of the Network Architecture and Design Team Task VI.

2.4.2 Caveats of Optimization Strategy

Several factors, some inherent in network design and some indicative of plan implementation, have a significant potential influence on the accuracy of network optimization efforts. The data used in this report is derived from a sample period of only one month which does not allow observation of peak months or growth. As discussed in Section 2.2, optimization is dependent on the data collected from router MIB statistics.

Incorrect or incomplete router configuration partially or completely nullifies extrapolations regarding network optimization. However, this caveat does not negate the optimization strategy but focuses attention on the need for **standards and management authority**.

The available link bandwidths were limited to the following: 64K, 128K, 256K, 512K, and T1. The reasons for bandwidth limitation were based on probable sizes needed by traffic demands, acceptable tariffs for these sizes, and the ratio of fractional T1 to a full T1 based on costing.

To insure that a particular application meets its network performance requirements, knowledge of the full routing path is required; guaranteeing that appropriate link capacity is available for the entire path. SNMP traffic analysis is limited to the source and destination router information. Task V (Application Level Traffic Study) of the initial TEN Design focuses on measurement of application traffic and performance. Finally, accurate knowledge of business requirements, applications, type and implementation, is critical to accurately minimize cost by optimizing link bandwidth. The current optimization recommendations are as complete as possible with the information now available.

2.4.3 Ramifications of TEN Project Optimization Strategy

The network link optimization strategy employed in the initial TEN Design will ultimately result, from a network utilization perspective, in a cost effective USDA network. Link capacity optimization does not solve other problems caused by the current network design. When additional network efforts (security, redundancy, and other management strategies) are implemented, the USDA network will better meet mission area business requirements. In addition to cost reduction, a less obvious but equally significant effect of network optimization is the improvement of business efficiency. More complete knowledge of network traffic improves the correlation between USDA business areas and the requirements of telecommunications interests. This improved correlation enhances the information transfer between business managers, applications managers, and network managers. It is implicit in the design strategy that optimization is dependent on the knowledge and accuracy of

network information. Constant monitoring of the network and use of better network information results in a continuous improvement of the USDA TEN.

2.5 Link Recapacitation Procedure

The procedure for network link recapacitation (Fig. 1) is highly dependent on the various features of the NetMaker XA[®] System. For this phase of optimization, the Capacitation Primitive of the Designer Tool. Link evaluation involves loading measured traffic data into Designer and allowing it to make link recapacitation recommendations based on optimization design criteria. Network design and evaluation criteria, discussed in Sections 2.3, 2.4. and the Addendum include a target usage of 50% but less than 70% of theoretical link capacity. By process iteration, NetMaker XA[®] recommends proper link size based on traffic, cost, and design parameters.

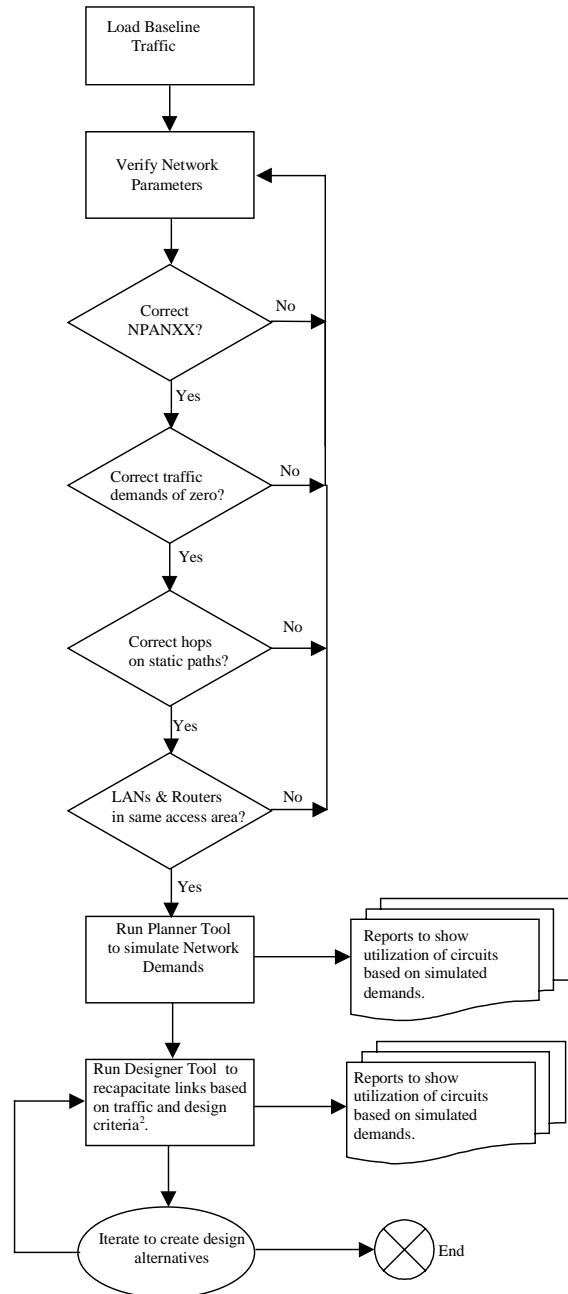


Figure 1 Network Recapacitation Procedure

3.0 Findings and Recommendations

3.1 Findings

3.1.1 General Department-wide Capacity Data

Based on the measured traffic demands and link accesses, 100% of the circuits of the USDA Data

Networks are incorrectly sized to some degree (Fig. 2). Of these incorrectly sized circuits, 84% are underutilized – candidates for link downsizing with appropriate cost savings.

Approximately 16% of the incorrectly sized links used in excess of 50% and require timely attention to avoid a mission performance impact (4% of the links exceed 70% requiring immediate attention).

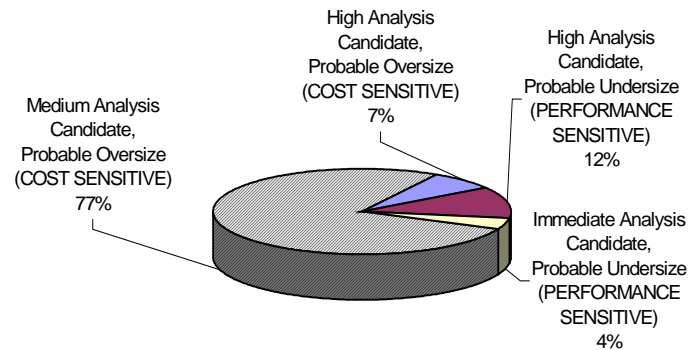


Figure 2 USDA Network Link Capacities

3.1.2 Capacity data of specific data networks

Table 2 provides statistics for the January 30, 1998 Discovered Network. WAN Access Cost is that portion of the monthly recurring service cost directly associated with the access into the Network. Frame Relay PVC or port charges are not included since the analysis focused on modifying the access links only. The Utilization statistics are based on the circuits which had traffic demands captured during the baseline process. For a significant number of circuits, there were no traffic demands captured.

WAN Access Cost ²	High Percent Utilization	Average Percent Utilization
\$469,829	25.5 %	5.4 %

Table 2 Measured Capacity and Utilization - Original Circuit Information

² WAN Access Cost = (number type DTS links)(cost/ DTS link) + (number FR links)(cost/ FR link)

3.2 Recommendations

3.2.1 Cost effect of implementing capacity recommendation

Table 3 presents the statistics for the recapacitated January 30, 1998 Discovered Network. WAN access cost is that portion of the monthly recurring service cost directly associated with the access into the Network.

WAN Access Cost ²	High Percent Utilization	Average Percent Utilization
\$360,580	43.0 %	9.4 %

Table 3 Link Capacity and Utilization - Recapacitated Circuit Information

If recapacitation recommendations are implemented, the total projected cost reduction for USDA is \$1.3 million per year; a relative cost savings of 23% . The relative cost saving better reflects the limitations noted in Section 2.4.2 and therefore is a better indicator of potential savings. It is also important to remember that this cost savings is the result only of phase one – recapacitation – of the network optimization strategy.

The analysis categories referenced in Section 2.3.1, Circuit Analysis Priority, reflect the general guidelines for sizing circuits. Because the baseline traffic has limitations, any further resizing using the capacitation primitive in NetMaker's Designer module is ineffective (this observation is from empirical evidence). Individual circuits may move from one category to another during the capacitation phase, but the number of circuits in a category remained approximately the same. This indicates that more specific traffic demands having more routing information is necessary to achieve an optimal design. This correlates well with the design process and the next step of application modeling.

Table 4 presents the WAN links access capacities for the initial network and the recapacitated network. These links produce the numbers in the 513-766K, 769-1024K, and 1025-1536K categories.

TELECOMMUNICATIONS ENTERPRISE NETWORK DESIGN

Development of Network Capacity Design – Task IV

Link Categories	Initial	Recapacitated
0-64K	225	307
65-128K	61	95
129-256K	52	193
385-512K	17	26
513-768K	35	3
769-1024K	4	1
1025-1536K	49	5
T1 ³	258	73
E1	1	0
Total	702	703

Table 4 Number of Links by Category Before and After Recapacitation

3.2.2 Performance effect of implementing capacity recommendations

Assuming that measured traffic demands accurately reflect the traffic between routers, there should be no noticeable decrease in performance for those circuits that are downsized. For the circuits which need to be upgraded, the performance should increase and benefit network customers.

4.0 Conclusions

A network link sizing metric to indicate the volume of link traffic is a complex function. When a link is underutilized, cost is the primary optimizing criterion since money is spent on unnecessary bandwidth. However, at peak periods of utilization the driving principle for link optimization is performance since an over-utilized link is operating at reduced performance. Consequently, a metric was developed that relates average usage and peak usage with a priority for link optimization analysis (Table 1). Application of the link optimization algorithm indicates that 100% of current USDA Data Networks links are incorrectly sized; 84% are underutilized (average utilization is less than 15% of link capacity) and 16% exceed 50% of link theoretical capacity (4% of these are in excess of 70%). Efficient link utilization needs better management oversight.

Seasonal variations in traffic, application performance requirements, and planned traffic increases were not considered when making recommendations. As a consequence, Agency recommendations (Addendum) are conservative. The recommendations should be used as a tool to improve utilization and performance of the existing network.

Link capacity optimization, the first phase of the overall optimization strategy for the USDA Data Networks is successfully approached using various the NetMaker XA[®]

³ One circuit was added at site requiring bandwidth greater than provided by a single T1.

TELECOMMUNICATIONS ENTERPRISE NETWORK DESIGN

Development of Network Capacity Design – Task IV

System tools. An optimization procedure using measured link data and iterative analysis provides a viable route to correctly sizing network links.

Using network link information collected for each Agency, link capacity was evaluated using the NetMaker XA System. The results demonstrate that, if all underutilized links are downsized and all over-utilized links are increased, the potential Department-wide cost savings is 23% while also improving the performance of 16% of the USDA links.

Addendum

NetMaker XA[®] Designer was used to recapacitate access links to routers. The following criteria and restrictions were used to evaluate the design of the network:

- Utilization was the determining factor in resizing recommendations with the high traffic demand driving the bandwidth recommendation. Seasonal variations in traffic, application performance requirements, and planned traffic increases were not considered when making these recommendations.
- Only access links to the WAN or Frame Relay POP from the router were candidates for resizing. No Frame Relay PVC's were resized.
- Existing links could not be deleted, but could be resized.
- Links with no traffic demands were not resized.
- For resizing, the set of available link bandwidths was limited to 64K, 128K, 256K, 512K, and T1.
- Cost for link access was based on tariffs contained in NetMaker database.
- Traffic demands were captured on weekdays from January 12, 1998 through February 6, 1998 from 4:00 AM through 7:00 PM (MST).